

Activated Random Walks*

Leonardo T. Rolla

Contents

1	Overview	2
1.1	The Activated Random Walk reaction-diffusion model	2
1.2	Infinite conservative system, fixation, phase transition	3
1.3	Physical motivation: “self-organized criticality”	3
1.4	Predictions	4
1.5	Difficulties	5
1.6	Results	5
2	Definitions and Diaconis-Fulton construction	8
2.1	Notation	8
2.2	Diaconis-Fulton construction	9
3	One-dimensional counting arguments	13
4	Exploring the instructions in advance	16
5	Higher-dimensional arguments	19
6	A multi-scale argument	24
7	Arguments using particle-wise constructions	27
7.1	Preliminaries	27
7.2	Results	29
	References	32

*Minicourse given at the workshop “Activated Random Walks, DLA, and related topics” at IMÉRA-Marseille, March 2015. **Preliminary version. July 15, 2015.**

1 Overview

1.1 The Activated Random Walk reaction-diffusion model

The Activated Random Walk model is defined as follows.

Particles sitting on the graph \mathbb{Z}^d can be in state A for *active* or S for *passive*. Each active particle, that is, each particle in the A state, performs a continuous-time random walk with jump rate $D_A = 1$ and with translation-invariant jump distribution given by a probability $p(\cdot)$ on \mathbb{Z}^d .

Several active particles can be at the same site, and they do not interact among themselves. When a particle is alone, it may become passive, a transition denoted by $A \rightarrow S$, which occurs at a *sleeping rate* $0 < \lambda \leq \infty$.

In other words, each particle carries two clocks, one for jumping and one for sleeping. Once a particle is passive, it stops moving, i.e., it has jump rate $D_S = 0$, and it remains passive until the instant when another particle is present at the same vertex. At such an instant the particle which is in S state flips to the A state, giving the transition $A + S \rightarrow 2A$.

If the clock rings for a particle to sleep while it shares a vertex with other particles, the tentative transformation $A \rightarrow S$ is overridden by the instantaneous reaction $A + S \rightarrow 2A$, so this attempt to sleep has no effect on the system state.

Jump distribution: $x \rightarrow x + z$ with probability $p(z)$.

Diffusion: jump at rates $D_A = 1$ and $D_S = 0$. No interaction.

Reactions: $A \rightarrow S$ at rate $0 < \lambda \leq \infty$, $A + S \rightarrow 2A$ at rate ∞ .

Notation: $\eta_t(x)$ denotes both the number and type of particles at site x at time t .

A particle in the S state stands still forever if no other particle ever visits the vertex where it is located. At the extreme case $\lambda = +\infty$, when a particle visits an empty site, it becomes passive instantaneously. This case is equivalent to *internal diffusion-limited aggregation* with infinitely many sources.

We have described *local rules* for the system to evolve. In order to fully describe the system, we need to specify on which subset of \mathbb{Z}^d this dynamics will occur, what are the boundary conditions, and the initial state at $t = 0$.

1.2 Infinite conservative system, fixation, phase transition

Consider a system running on the whole graph \mathbb{Z}^d , and such that the initial configuration η_0 is i.i.d. Poisson with parameter μ .

A phase transition in this systems arises from a conflict between a *spread of activity* and a *tendency for this activity to die out*, and the transition point separates an active and an absorbing phase in which the dynamics gets eventually extinct in any finite region.

We say that the system *locally fixates* if $\eta_t(x)$ is eventually constant for each x , otherwise we say that the system *stays active*.

It can be shown using ergodicity that the probability that the system fixates is either 0 or 1. Moreover,

If the system fixates for given λ and μ , then it fixates for larger λ and smaller μ .

As a consequence,

$$\mathbb{P}(\text{fixation}) = \begin{cases} 1, & \mu < \mu_c, \\ 0, & \mu > \mu_c, \end{cases}$$

where the critical density $\mu_c = \mu_c(\lambda) \in [0, \infty]$ is non-decreasing in $\lambda \in [0, \infty]$.

1.3 Physical motivation: “self-organized criticality”

In contrast with the above description, consider now the following evolution.

Driven-dissipative system. For large L , let the system run on the finite box $V = [0, L]^d$. New particles are added to the bulk of V at constant rate. When a particle hits the boundary of V it is killed. The reaction-diffusion dynamics is run at a much faster speed, so that the whole box is stabilized between two arrivals. We then let $t \rightarrow \infty$ to reach stationarity, then take $L \gg 1$ to have a state on \mathbb{Z}^d .

The relation between self-organized and ordinary criticality is understood as follows.

On the one hand, “self-organized criticality” appears in the parameter-free, driven-dissipative evolution described above. In this dynamics, when the average density μ inside the box is too small, mass tends to accumulate. When it is too large, there is intense activity and a substantial number of particles is absorbed at the boundary. With

this *carefully designed* mechanism, the model is attracted to a *critical state* with an average density given by $0 < \mu_c < \infty$, though it was not explicitly tuned to this critical value.

On the other hand, the corresponding conservative system in infinite volume exhibits ordinary criticality in the sense that its dynamics fixate for $\mu < \mu_c$ and do not fixate for $\mu > \mu_c$, and moreover the *critical exponents* of the finite-volume addition-relaxation dynamics are related to those of the conservative dynamics in infinite volume.

1.4 Predictions

The behavior of the ARW is expected to be the following.

The critical density satisfies $\mu_c \rightarrow 0$ as $\lambda \rightarrow 0$ and $\mu_c \rightarrow 1$ as $\lambda \rightarrow \infty$. See Figure 1. The value of μ_c should not depend on the particular μ -parametrized distribution of the initial configuration (geometric, Poisson, etc.).

At $\mu = \mu_c$, the density of active particles vanishes as $t \rightarrow \infty$, but we conjecture that the system does not fixate in this case.

The asymptotic decay of density of activity as $t \gg 1$ when $\mu = \mu_c$ should obey a power law. Also, for the stationary regime, i.e., letting $t \rightarrow \infty$ first, the density of activity should decay with a power law as $0 < \mu - \mu_c \ll 1$. More predictions can be found at [DRS10].




Figure 1: Predictions

1.5 Difficulties

For the ARW, we would like to describe the critical behavior, the scaling relations and critical exponents, and whether the critical density is the same as the long-time limit attained in the driven-dissipative version. These questions are however far beyond the reach of current techniques.

The first apparent difficulty lies in the fact that this system is not attractive (i.e., its evolution does not preserve monotonicity of configurations). This is overcome by considering a Diaconis-Fulton kind of construction, or an explicit construction in terms of a collection of random walks, rather than the classical Harris graphical construction. These constructions allow for different kinds of arguments which have proven to be very useful.

However, another particularity of this model brings severe difficulties, which is particle conservation. In particular, this seems to rule out any energy-entropy kind of argument. These arguments typically go as follows. One first identifies some structure that is intrinsic to the occurrence of events that conjecturally should not occur. The number of possible structures is then estimated, and when it is overwhelmed by the high energy cost needed to construct them, one can show that such events have vanishing probability. This approach has proven successful in perturbative statistical mechanics. However, the conservation of particles in the system gives rise to intricate long-range effects, which makes it difficult to find a suitable structure behind self-sustained activity for large periods of time.

Lack of attractiveness. Overcome by using other constructions rather than Harris’.

Conservation of particles. Rules out any “energy x entropy” kind of approach.

1.6 Results

We end this overview by recalling the existing results in the literature about the phase transition of the ARW on the infinite lattice \mathbb{Z}^d . In the next sections we should go through the proof of all the results mentioned here.

The results will be presented in the same framework, by making the following assumptions.

We shall assume that: (i) the initial distribution is i.i.d. Poisson parametrized by μ , and (ii) jumps are to nearest neighbors, that is, $p(z) = 0$ unless $\|z\| = 1$.

For each result, the assumption of i.i.d. initial condition and a Poisson law for its marginals can be more or less relevant depending on the proof. We cannot make any other assumptions about the jump distribution, since some of the existing results have been proved only for biased jumps, whereas some others only in the complementary case.

Each result can be generalized in different directions. The reader interested in minimal hypotheses is referred to the original articles.

One-dimensional directed walks

For the particular situation when particles only jump to their right, all the above predictions about the phase transition curve can be proved. More precisely,

$$\mu_c = \frac{\lambda}{1 + \lambda},$$

and there is not fixation at $\mu = \mu_c$. This appeared in [CRS14] although it was known before. Yet, much remains to be understood about its critical behavior. In particular, the scaling limit of the flow process has been studied in [CRS14] for the extreme case $\lambda = \infty$, but remains open elsewhere on the critical curve.




Figure 2: Results for $d = 1$, totally asymmetric jumps

One-dimensional systems

In studying a different model, arguments found in [KS06] imply that for high enough μ the ARW with $\lambda = \infty$ model stays active. In terms of the phase diagram, $\mu_c < \infty$ for

any $\lambda \leq \infty$. It was shown in [RS12] that there is no fixation at $\mu = 1$, so in particular $\mu_c \leq 1$ for any $\lambda \leq \infty$.

Finally, in [Tag14] it was shown that when the jump distributions is biased, $\mu_c < 1$ for any $\lambda < \infty$, and $\mu_c \rightarrow 0$ as $\lambda \rightarrow 0$. It should be noted that the lower bound obtained for μ_c tends to 1 if the bias is small, for any $\lambda > 0$, and that the case of unbiased jumps remains open.

The first result on fixation in this context appeared in [RS12], showing that

$$\mu_c \geq \frac{\lambda}{1 + \lambda}.$$

In terms of the phase diagram, $\mu_c > 0$ for all λ , and $\mu_c \rightarrow 1$ as $\lambda \rightarrow \infty$.




Figure 3: Results for $d = 1$, general jumps

Two and more dimensions

In [KS06], the fact that $\mu_c < \infty$ was proved for any dimension. It was then shown that $\mu_c \leq 1$ by [She10] and [AGG10], using different methods. This was extended in [CRS14], who showed that there is no fixation at $\mu = 1$.

In [Tag14], it was shown that $\mu_c < 1$ for sufficiently biased jump distributions and some $\lambda > 0$. Unlike the one-dimensional case, the proof gives neither $\mu_c < 1$ for all $\lambda < \infty$, nor $\mu_c \rightarrow 0$ as $\lambda \rightarrow 0$. The case of small or zero bias remains open.

The first result on fixation for arbitrary dimensions appeared in [She10], who shows that $\mu_c > 0$ when $\lambda = \infty$. This was extended in [CRS14] who show that $\mu_c = 1$ when $\lambda = \infty$.

Finally, in [ST14], it was shown that when the jump distribution is unbiased, $\mu_c > 0$ for any $\lambda > 0$. The methods of [ST14] do not allow to prove that $\mu_c \rightarrow 1$ as $\lambda \rightarrow \infty$, though. The case of biased jumps and $\lambda < \infty$ remains widely open.




Figure 4: Results for $d \geq 2$

2 Definitions and Diaconis-Fulton construction

2.1 Notation

Let $\mathbb{N}_0 = \{0, 1, 2, \dots\}$ and $\mathbb{N}_\varrho = \mathbb{N}_0 \cup \{\varrho\}$. The state of the ARW at time $t \geq 0$ is given by $\eta_t \in \Sigma = (\mathbb{N}_\varrho)^{\mathbb{Z}^d}$. In this setting, $\eta_t(x) = n$ means that, at time t , site x contains n active particles if $n \in \mathbb{N}_0$, or one passive particle if $n = \varrho$.

We turn \mathbb{N}_ϱ into an ordered set by letting $0 < \varrho < 1 < 2 < \dots$. We also let $|\varrho| = 1$, so $|\eta_t(x)|$ counts the number of particles regardless of their state. We define $\llbracket n \rrbracket$ to be n if $n \geq 1$, and 0 if $n = 0$ or ϱ , so $\llbracket \eta_t(x) \rrbracket$ counts the number of active particles.

To add a particle to a site we define $\varrho + 1 = 2$, which represents the $A + S \rightarrow 2A$ transition. We also define $1 \cdot \varrho = \varrho$ and $n \cdot \varrho = n$ for $n \geq 2$, which represent the transitions $A \rightarrow S$ and $2A \rightarrow A + S \rightarrow 2A$, respectively.

It will be convenient to consider as *acceptable* the operations $\varrho - 1 = 0$ and $\varrho \cdot \varrho = \varrho$, although they do not appear in the dynamics. The operations $0 \cdot \varrho$ and $0 - 1$, on the other hand, are not acceptable.

The process evolves as follows. For each site x , a Poisson clock rings at rate $(1 + \lambda)\llbracket \eta_t(x) \rrbracket$. When this clock rings, the system goes through the transition $\eta \rightarrow \tau_{x\varrho}\eta$ with probability

$\frac{\lambda}{1+\lambda}$, otherwise $\eta \rightarrow \tau_{xy}\eta$ with probability $p(y-x)\frac{1}{1+\lambda}$. The transitions are given by

$$\tau_{xy}\eta(z) = \begin{cases} \eta(x) - 1, & z = x \\ \eta(y) + 1, & z = y \\ \eta(z), & \text{otherwise,} \end{cases} \quad \text{and} \quad \tau_{x\varrho}\eta(z) = \begin{cases} \varrho \cdot \eta(x), & z = x \\ \eta(z), & \text{otherwise.} \end{cases}$$

We assume that $\eta_0(x) \in \mathbb{N}_0$ for all x a.s., and use \mathbb{P}^ν denote the law of $(\eta_t)_{t \geq 0}$, where ν denotes the distribution of η_0 .

2.2 Diaconis-Fulton construction

The Diaconis-Fulton representation enables us to exploit the combinatorial nature of fixation. Due to particle exchangeability, this representation extracts precisely the part of the randomness that is relevant for the phase transition, focusing on the total number of jumps and leaving aside the order in which they take place. It is suitable for studying path traces, total occupation times, and final particle positions, but precludes the analysis of quantities for which the order of the jumps does matter, such as correlation functions or local shape properties.

We summarize the main points of this section for later reference.

- *Toppling* x is an operator that reduces the state at x and increases elsewhere.
- *Legal topplings* are those actually performed by the system.
- The order of topplings is irrelevant and we are free to choose it.
- *Acceptable toppling* is an artificial operation that extends the definition of legal topplings, and provides *upper bounds* for activity.
- On the other hand, legal topplings provide *lower bounds* for activity.
- This construction determines fixation or non-fixation for the continuous-time model.
- There is a zero-one law for fixation.

In this section, η denotes a configuration, and we do not deal with a continuous-time evolution anymore. We say that site x is *unstable for the configuration* η if $\eta(x) \geq 1$. Unstable sites can *topple*, by applying τ_{xy} or $\tau_{x\varrho}$ to η . Toppling an unstable site is *legal*. If $\eta(x) = \varrho$, we still say that toppling x is *acceptable*. A legal toppling is also acceptable. Consider a field of instructions $\mathcal{I} = (\tau^{x,j})_{x \in \mathbb{Z}^d, j \in \mathbb{N}}$. Later on we will choose \mathcal{I} random, but for now \mathcal{I} denotes a field that is fixed, and η denotes any configuration.

Let $h = (h(x); x \in \mathbb{Z}^d)$ count the number of topplings at each site, usually started from $h \equiv 0$. The toppling operation at x is defined by

$$\Phi_x(\eta, h) = (\tau^{x, h(x)+1}\eta, h + \delta_x),$$

and $\Phi_x\eta$ is a short for $\Phi_x(\eta, 0)$. Again, Φ_x is *legal* for η if $\eta(x) \geq 1$, and *acceptable* for η if $\eta(x) \geq \varrho$.

Sequences of topplings and local properties

Let $\alpha = (x_1, \dots, x_k)$ denote a finite sequence of sites. We define $\Phi_\alpha = \Phi_{x_k} \Phi_{x_{k-1}} \cdots \Phi_{x_1}$. We say that Φ_α is *legal* (respectively *acceptable*) for η if Φ_{x_l} is legal (respectively acceptable) for $\Phi_{(x_1, \dots, x_{l-1})}\eta$ for each $l = 1, \dots, k$. In this case we say that α is a *legal sequence* (respectively *acceptable sequence*) of topplings for η .

Let $m_\alpha = (m_\alpha(x); x \in \mathbb{Z}^d)$ be given by $m_\alpha(x) = \sum_l \mathbb{1}_{x_l=x}$, the number of times the site x appears in α . We write $m_\alpha \geq m_\beta$ if $m_\alpha(x) \geq m_\beta(x) \forall x$, and $\tilde{\eta} \geq \eta$ if $\tilde{\eta}(x) \geq \eta(x) \forall x$. We also write $(\tilde{\eta}, \tilde{h}) \geq (\eta, h)$ if $\tilde{\eta} \geq \eta$ and $\tilde{h} = h$.

Let x be a site in \mathbb{Z}^d and $\eta, \tilde{\eta}$ be configurations.

- 1) Local Abelianness.** If α and β are acceptable sequences of topplings for the configuration η , such that $m_\alpha = m_\beta$, then $\Phi_\alpha\eta = \Phi_\beta\eta$.
- 2) Mass comes from outside.** If α and β are acceptable sequences of topplings for η such that $m_\alpha(x) \leq m_\beta(x)$ and $m_\alpha(z) \geq m_\beta(z)$ for $z \neq x$, then $\Phi_\alpha\eta(x) \geq \Phi_\beta\eta(x)$.
- 3) Monotonicity of stability.** If site x is unstable for the configuration η , and if $\tilde{\eta}(x) \geq \eta(x)$, then x is also unstable for the configuration $\tilde{\eta}$.
- 4) Monotonicity of topplings.** If $\tilde{\eta} \geq \eta$ and Φ_x is legal for η , then Φ_x is legal for $\tilde{\eta}$ and $\Phi_x\tilde{\eta} \geq \Phi_x\eta$.

Proof. The two last properties are immediate. Suppose that $m_\alpha = m_\beta$.

For convenience, define the operators $n \oplus = n + 1$ on \mathbb{N}_ϱ , we well as $n \ominus = n - 1$ and $n \odot = n \cdot \varrho$ on $\mathbb{N}_\varrho \setminus \{0\}$. With this notation, whenever $n \ominus$ is acceptable (i.e., $n \neq 0$) we have $n \ominus \oplus = n \oplus \ominus$. Analogously, whenever $n \odot$ is acceptable (i.e., $n \neq 0$) we have $n \odot \oplus = n \oplus \odot$. Therefore, within any acceptable sequence of operations, replacing $\ominus \oplus$ and $\odot \oplus$ by $\oplus \ominus$ and $\oplus \odot$ yields an acceptable sequence with the same final outcome.

Notice that $\Phi_\alpha\eta(x)$ is given by $\eta(x)$ followed by a number of \oplus , \ominus and \odot 's. The number of times each operator appears is determined by \mathcal{I} and m_α and is thus the same for

$\Phi_\beta\eta(x)$. Their actual order depends on the sequence, but the internal order of the \ominus and \odot operators is determined by $(\tau^{x,j})_j$ and is the same for both $\Phi_\alpha\eta(x)$ and $\Phi_\beta\eta(x)$. As a consequence, we can apply the above identities to move the \oplus 's to the left, yielding then identical sequences for $\Phi_\alpha\eta(x)$ and $\Phi_\beta\eta(x)$, which proves the first property.

The second property follows from a similar argument. Suppose $m_\alpha(x) \leq m_\beta(x)$ and $m_\alpha(z) \geq m_\beta(z)$ for $z \neq x$. Again $\Phi_\alpha\eta(x)$ is given by $\eta(x)$ followed by a number of \oplus , \ominus and \odot 's. The number of times that operator \oplus appears depends on $m_\alpha(z)$, $z \neq x$, and is bigger than for $\Phi_\beta\eta(x)$, and the number of times that operators \ominus and \odot appear depend on $m_\alpha(x)$, and is smaller than for $\Phi_\beta\eta(x)$. Pushing the \oplus 's to the left we get that $\Phi_\alpha\eta(x)$ is written in the same way as $\Phi_\beta\eta(x)$, perhaps with a few extra \oplus 's in the beginning, and a few missing \ominus and \odot 's in the end, so $\Phi_\alpha\eta(x) \geq \Phi_\beta\eta(x)$. \square

Global properties

What follows is valid for any model satisfying the above four properties.

Let V denote a finite subset of \mathbb{Z}^d . We say that η is *stable in V* if every $x \in V$ is stable for η . We say that α is *contained in V* , and write $\alpha \subseteq V$, if every x appearing in α is an element of V . We say that α *stabilizes η in V* if α is acceptable for η and $\Phi_\alpha\eta$ is stable in V .

Least Action Principle. If α is an acceptable sequence of topplings that stabilizes η in V , and $\beta \subseteq V$ is a sequence of topplings that is legal for η , then $m_\beta \leq m_\alpha$.

Proof. Let $\beta \subseteq V$ be legal and $m_\alpha \not\leq m_\beta$. Write $\beta = (x_1, \dots, x_k)$ and $\beta^{(j)} = (x_1, \dots, x_j)$ for $j \leq k$. Let $\ell = \max\{j : m_{\beta^{(j)}} \leq m_\alpha\} < k$ and $y = x_{\ell+1} \in V$. Since β is legal, y is unstable in $\Phi_{\beta^{(\ell)}}\eta$. But $m_{\beta^{(\ell)}} \leq m_\alpha$ and $m_{\beta^{(\ell)}}(y) = m_\alpha(y)$. By the Properties 2 and 3, y is unstable for $\Phi_\alpha\eta$ and therefore α does not stabilize η in V . \square

For each finite set V and configuration η , we define

$$m_{V,\eta} = \sup_{\beta \subseteq V \text{ legal}} m_\beta.$$

The Least Action Principle says that

$$m_\beta \leq m_{V,\eta} \leq m_\alpha$$

for any legal sequence β contained in V and any acceptable sequence α stabilizing η in V . This provides a very good source of lower and upper bounds for $m_{V,\eta}$.

Global Abelianness. If α and β are both legal toppling sequences for η that are contained in V and stabilize η in V , then $m_\alpha = m_\beta = m_{V,\eta}$. In particular, $\Phi_\alpha \eta = \Phi_\beta \eta$.

Proof. Apply the Least Action Principle in two directions: $m_\beta \leq m_{V,\eta} \leq m_\alpha \leq m_\beta$. \square

We say that η is *stabilizable in V* if there exists an acceptable sequence α that stabilizes η in V . Notice that this provides finite upper bounds for legal sequences $\beta \subseteq V$, which in turn implies the existence of a sequence contained in V that is both legal and stabilizing. Indeed, if one tries to perform legal topplings in V indefinitely, on the one hand one has to eventually stop since there is a finite upper bound, and on the other hand this can only be stopped if there are no more unstable sites.

Monotonicity. If $V \subseteq \tilde{V}$ and $\eta \leq \tilde{\eta}$, then $m_{V,\eta} \leq m_{\tilde{V},\tilde{\eta}}$.

Proof. Let $\beta \subseteq V$ be legal for η . By successively applying Properties 3 and 4, β is also legal for $\tilde{\eta}$. Since $\beta \subseteq \tilde{V}$ in this case, the inequality follows from the definition of $m_{V,\eta}$. \square

By monotonicity, the limit

$$m_\eta = \lim_{V \uparrow \mathbb{Z}^d} m_{V,\eta}$$

exists and does not depend on the particular sequence $V \uparrow \mathbb{Z}^d$ (the limit is given by the supremum over finite V). A configuration η is said to be *stabilizable* if $m_\eta(x) < \infty$ for every $x \in \mathbb{Z}^d$.

Fixation for the stochastic dynamics

Recall that \mathbb{P}^ν denotes the law of the process $(\eta_t)_{t \geq 0}$ with values on $\mathbb{N}_\varrho^{\mathbb{Z}^d}$. Assume that distribution ν of η_0 on $\mathbb{N}_0^{\mathbb{Z}^d}$ is ergodic and has finite density $\nu(\eta(\mathbf{0})) < \infty$. We further assume that the support of $p(\cdot)$ generates \mathbb{Z}^d and not a sublattice.

To avoid extra notation, we define the field \mathcal{I} on the same probability space \mathbb{P}^ν , independently of η_0 , and distributed as follows. For each $x \in \mathbb{Z}^d$ and $j \in \mathbb{N}$, choose $\tau^{x,j}$ as τ_{xy} with probability $\frac{p(y-x)}{1+\lambda}$ or $\tau_{x\varrho}$ with probability $\frac{\lambda}{1+\lambda}$, independently over x and j .

Stochastic Dynamics. $\mathbb{P}^\nu(\text{fixation of } (\eta_t)_{t \geq 0}) = \mathbb{P}^\nu(m_{\eta_0}(\mathbf{0}) < \infty) = 0 \text{ or } 1$.

For the proof, the reader is referred to [RS12, v2 on arXiv].

3 One-dimensional counting arguments

[CRS14] ($d = 1$. Directed walks). $\mu_c = \frac{\lambda}{1+\lambda}$. No fixation for $\mu = \mu_c$.

Proof. Fixation is equivalent to $m_{\eta_0}(\mathbf{0}) < \infty$ almost surely. This is equivalent to $m_{V_L, \eta_0}(\mathbf{0})$ being tight as $L \rightarrow \infty$, where $V_L = \{-L, \dots, 0\}$. The latter is equivalent to tightness of the number of particles which jump into $\mathbf{0}$ when V_L is stabilized with legal topplings. See Figure 5.




Figure 5: The set V_L , part of the configuration η , and arrow indicating the unknown number of particles which jump into $\mathbf{0}$ when stabilizing $[-L, 0]$

Let us stabilize V_L by exhausting each site, from left to right. More precisely, topple site $x = -L$ until each of the $\eta_0(-L)$ particles either moves to $x = -L + 1$ or sleeps, and let Y_0 be indicator of the event that the last particle remained sleeping on $x = -L$. Conditioned on $\eta_0(-L)$, the distribution of Y_0 is Bernoulli with parameter $\frac{\lambda}{1+\lambda}$ (in case $\eta_0(-L) = 0$, sample Y_0 independently of anything else). The number of particles which jump from $x = -L$ to $x = -L + 1$ is given by $N_1 := [\eta_0(-L) - Y_0]^+$. See Figure 6.




Figure 6: Stabilizing $[-L, 0]$ from left to right. Disks represent particles initially present at each site, and circles represent particles that arrive from the left. Red boxes indicate particles which became passive when they were left alone.

Note that, after stabilizing $x = -L$, there are $N_1 + \eta_0(-L + 1)$ particles at $x = -L + 1$. Let the site $x = -L + 1$ topple until it is stable, and denote by Y_1 the indicator of the event that the last particle remained passive on $x = -L + 1$. The number of particles which jump from $x = -L + 1$ to $x = -L + 2$ is given by $N_2 := [N_1 + \eta_0(-L + 1) - Y_1]^+$. By iterating this procedure, the number N_{i+1} of particles which jump from $x = -L + i$ into $x = -L + i + 1$ after stabilizing $x = -L, -L + 1, \dots, -L + i$ is given by

$$N_{i+1} = [N_i + \eta_0(-L + i) - Y_i]^+,$$

where $N_0 = 0$. The number of particles which jump into $\mathbf{0}$ while stabilizing V_L equals N_L . Note that the sequence $(N_i)_{i=0,1,\dots,L}$ is a random walk on \mathbb{N}_0 , with independent increments distributed as $\eta_0(x) - Y$, reflected at 0. So the relevant quantity is

$$\mathbb{E}[\eta_0(-L + k) - Y_k] = \mu - \frac{\lambda}{1 + \lambda}.$$

If $\mu < \frac{\lambda}{1+\lambda}$, the walk is positive recurrent, which implies tightness of N_i , and thus fixation. If $\mu > \frac{\lambda}{1+\lambda}$, the walk is transient, and $\mathbb{P}[N_L \geq \frac{1}{2}(\mu - \frac{\lambda}{1+\lambda})L]$ is large for large L , so there is no fixation. If $\mu = \frac{\lambda}{1+\lambda}$, the walk is null-recurrent, so there is no tightness and thus no fixation. \square

($d = 1$). No fixation when $\mu = 1$.

Proof. Let $\mu = 1$. By the CLT, the probability that η_0 contains at least $L + 2\sqrt{L}$ particles on $V_L = [0, L]$ is at least $2\delta > 0$, uniformly on L . On this event, at least $2\sqrt{L}$ particles will visit either $x = \mathbf{0}$ or $x = L$ when we stabilize $[0, L]$. Therefore,

$$2\delta \leq \mathbb{P}(m_{V_L, \eta_0}(\mathbf{0}) \geq \sqrt{L}) + \mathbb{P}(m_{V_L, \eta_0}(L) \geq \sqrt{L}) \leq 2\mathbb{P}(m_{\eta_0}(\mathbf{0}) \geq \sqrt{L}).$$

Since this is true for any L , there is no fixation. \square

[Tag14] ($d = 1$, Biased walks). For any $\lambda < \infty$, $\mu_c < 1$. Moreover, $\lim_{\lambda \rightarrow 0} \mu_c(\lambda) = 0$.

Proof. Let $F(\lambda)$ be the probability that a walk which jumps at rate 1 and sleeps at rate λ , but only sleeps at $(-\infty, \mathbf{0}]$, is able to reach $+\infty$ without sleeping. A more explicit definition is $F(\lambda) = \mathbb{E}[(\frac{1}{1+\lambda})^T]$, where T counts how many steps a discrete-time random walk starting at $x = \mathbf{0}$ is found on $(-\infty, 0]$. We will show that $\mu_c \leq 1 - F(\lambda)$.

Suppose $\mu > 1 - F(\lambda)$. We will perform a sequence of legal toppings on $[-2L, 0]$ as follows.

In the first stage, we move each particle initially found on $[-L, -1]$ until it is alone, in order to have a configuration on $\{0, 1\}^L$. If particles reach $\{0\}$ they are stopped, but on $[-2L, -L]$ we still move them. With high probability, none of the biased random walks will go further to the left than $\log^2 L$. Therefore, after this step at most $\log^2 L$ particles

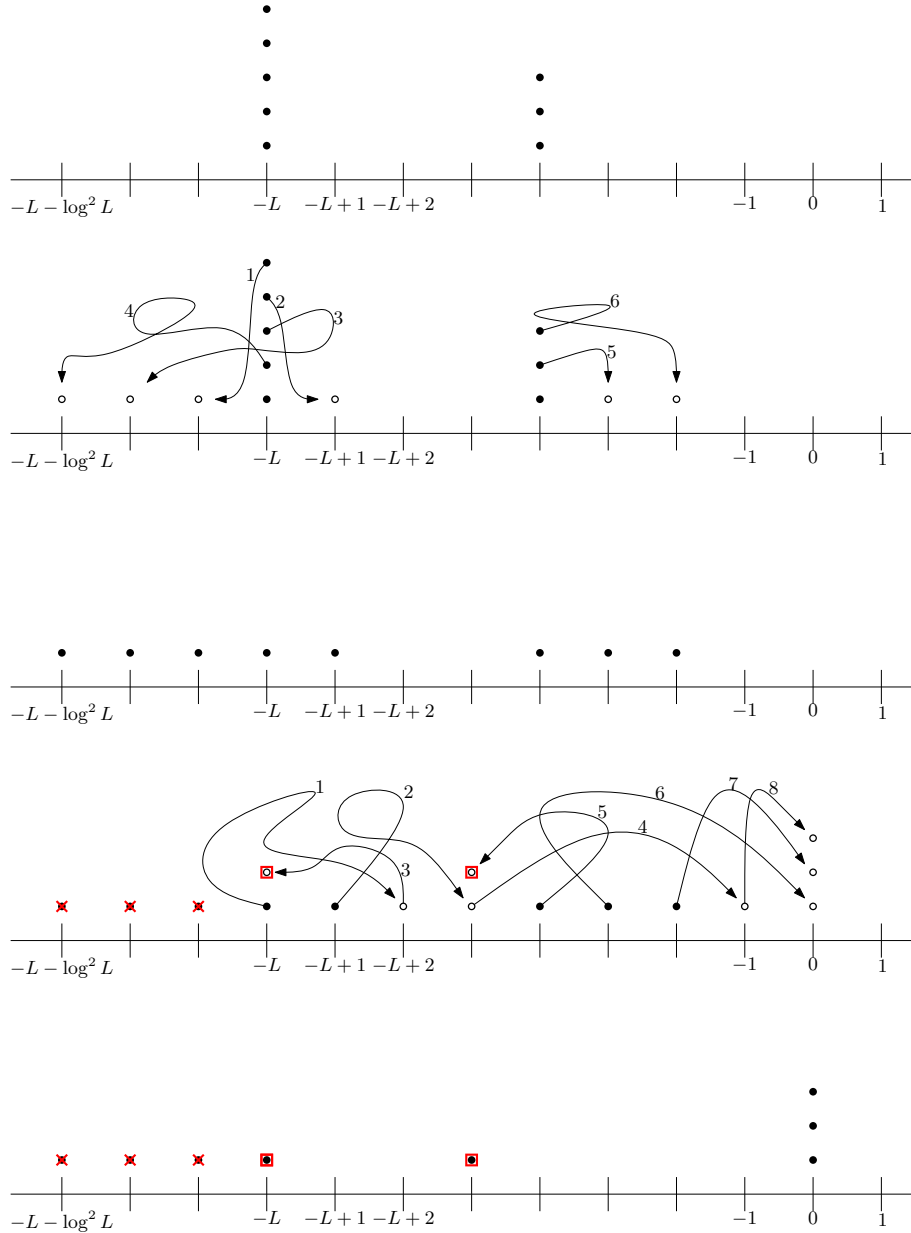


Figure 7: The two stages in the toppling procedure. After the second stage, a number of particles is found at 0 .

will be outside $[-L, 0]$.

We now move to the second stage. Let N_0 denote the number of particles on $A_0 = [-L, 0]$ after this first step. Notice that

At this point, all sites in this set, except site $\mathbf{0}$, have either 0 or 1 particle (*)

Let $x = -L$ and $A_1 = [-L + 1, 0]$. If x has no particles, then at A_1 there are $N_1 = N_0$ particles, and A_1 satisfies (*). If x has one particle, move this particle until it either sleeps on $[-\infty, x]$, or finds an empty site on $[x + 1, -1]$, or else it reaches $\{0\}$. At the end of this step, the number of particles present at A_1 will be either N_0 or $N_0 - 1$. The latter happens with probability at most $1 - F(\lambda)$. In any case the configuration on A_1 satisfies (*).

Move to $x = -L + 1$, define $A_2 = [-L + 2, 0]$, again move the particle found at x (if any) until it reaches an empty site. Again the number N_2 of particles in A_2 after this step is $N_1 - 1$ with probability at most $1 - F(\lambda)$, otherwise is equal to N_1 , and A_2 satisfies (*). Then move to $x = -L + 2$, define A_3 , N_3 , and so on.

Finally, when $x = -1$ we have the set $A_L = \{0\}$ containing N_L particles. By applying the LLN, both to the number of particles initially found on $[-L, 0]$, and to the number of particles lost on the second stage, we have that $N_L > (\mu - \epsilon)L - \log^2 L - [1 - F(\lambda) + \epsilon]L \gg 1$ with high probability, finishing the proof. \square

Notice that the upper bound for μ_c degenerates as the jumps become less biased.

4 Exploring the instructions in advance

[RS12] ($d = 1$). For any $\lambda > 0$, $\mu_c \geq \frac{\lambda}{1+\lambda}$.

Sketch of the proof. The proof uses an algorithm that tries to stabilize all the particles initially present in η_0 , following the instructions in \mathcal{I} , with the aid of acceptable topplings.

After describing the algorithm, we will show that, whenever it is successful, it implies that $m_{\eta_0}(\mathbf{0}) = 0$. We finally show that the algorithm is successful with positive probability if $\mu < \frac{\lambda}{1+\lambda}$. By the 0-1 law, this implies almost sure fixation of the ARW model.

I. Description of the algorithm.

The algorithm consists in applying a trapping procedure to each particle. This procedure explores \mathcal{I} until it identifies a suitable trap (given by a carefully chosen sleep instruction)

for the particle. To do that, it follows the path that the particle would perform if we always toppled the site it occupies, and stops when the trap has been chosen. In the absence of a suitable trap, the algorithm fails.

Remark that some of the explored instructions are actually not going to be used by the corresponding particle by the time it settles at the trap, leaving some *corrupted sites* that could interfere with the statistics of the subsequent steps.

If there is a particle at $\mathbf{0}$, we declare the procedure unsuccessful and stop. Otherwise, label the initial positions of the particles on \mathbb{Z} by $\cdots \leq x_{-3} \leq x_{-2} \leq x_{-1} < 0 < x_1 \leq x_2 \leq x_3 \leq \cdots$. Let $a_0 = 0$.

We now describe the trapping procedure for particle x_k , $k \geq 1$. Suppose that the first $k-1$ traps have been successfully set up at positions $0 < a_1 < a_2 < \cdots < a_{k-2} < a_{k-1} < x_{k-1}$. The set $[0, a_{k-1}]$ contains all traps and corrupted sites found so far.

The settling procedure starts with an *exploration*. Starting at x_k , examine and follow the instructions in \mathcal{I} one by one (whenever a sleep instruction is found, the next instruction at the same site is to be examined). Follow this exploration until reaching a_{k-1} .

Next we *set up the trap*. During the k -th exploration, we are sure to visit every $x \in B_k = [a_{k-1} + 1, x_k - 1]$. Moreover, the last instruction explored at each $x \in B_k$ is a jump to the left, see Figures 8 and 9. For each $x \in B_k$, the second last instruction may be a sleep instruction. If this is not the case for any $x \in B_k$, we declare the procedure unsuccessful and stop. Otherwise, let a_k be the leftmost site at which the second last instruction explored was a sleep instruction, and call this second last instruction the *k-th trap*.

Since the trap is a sleep instruction found immediately before the last instruction, which is a jump to the left, we are sure that the exploration path has not gone to the right of a_k after finding the trap, so *all the corrupted sites will be in $[a_{k-1} + 1, a_k]$* , see Figure 9. This procedure can be carried on indefinitely, as long as *all the steps are successful*.

II. If the algorithm is successful, then $m_{\eta_0}(\mathbf{0}) = 0$.

We will show that, following the instructions of \mathcal{I} , η_0 is stabilized in $V_n = [x_{-n}, x_n]$ with finitely many acceptable topplings, without toppling $\mathbf{0}$.

Let us first stabilize the particle that starts at x_1 . To that end, we successively topple the sites found by the first explorer, until it reaches the trap at a_1 . At this moment the particle will become passive, and the site a_1 will be stable. Notice that this is acceptable since the particle is following the same path that the explorer did, even if it will sometimes be passive.

Notice that, after the last visit to a_1 , the explorer did not go further to the right, so when




Figure 8: First exploration path. It starts at position x_1 of the first particle and stops when it reaches the origin. The horizontal axis represents the lattice, and above each site x there is a sequence of instructions $(\tau^{x,j})_j$. The bold arrows indicate the last jump found at each site $x \in [1, x_1 - 1]$, and the bold cross indicates a sleep instruction found just before the last jump, this being the leftmost such cross, whose location defines a_1 .

settling the first particle we use all the instructions examined so far, except some lying in $[a_0 + 1, a_1]$. Therefore, the same procedure can be applied to the second particle, as it will find the same instructions that determined the second exploration path.

Notice also that the first particle does not visit $\mathbf{0}$, and the second particle neither visits $\mathbf{0}$ nor a_1 , thus it is settled without activating the first particle. Following the same procedure, the k -th particle is settled at a_k , without ever visiting $\{0, a_1, a_2, \dots, a_{k-1}\}$, for all $k = 1, \dots, n$. After settling the n first particles in \mathbb{Z}_+ , we perform the analogous procedure for the first n particles in \mathbb{Z}_- .

This means that η_0 can be stabilized in V_n with finitely many acceptable topplings, not necessarily in V_n , and never toppling the origin. By the Least Action Principle, $m_{V_n, \eta}(\mathbf{0}) = 0$. Since it holds for all $n \in \mathbb{N}$ and $V_n \uparrow \mathbb{Z}$ as $n \rightarrow \infty$, this gives $m_\eta(\mathbf{0}) = 0$.




Figure 9: Second exploration path. It starts at position x_2 of the second particle and stops when it reaches a_1 . The regions in gray indicate the instructions already examined by the first explorer. The dark gray contains instructions examined but not used, whose locations determine the set of corrupted sites.

III. The algorithm is successful with positive probability.

For each site $x \in B_1$, the probability of finding a sleep instruction just before its last jump equals $\frac{\lambda}{1+\lambda}$, and this happens independently of the path and independently for each site. Thus, $a_1 - a_0$ is a random variable distributed as a geometric with parameter $\frac{\lambda}{1+\lambda}$, truncated at $x_1 - a_0$.

Since no corrupted sites were left outside $[a_0 + 1, a_1]$ in the first exploration, the inter-distance $a_2 - a_1$ is independent of a_1 . Moreover, its distribution is also geometric with parameter $\frac{\lambda}{1+\lambda}$. The same holds for $a_3 - a_2$, $a_4 - a_3$, etc. By the law of large numbers, $a_n \sim \frac{1+\lambda}{\lambda}n$. On the other hand, $x_n \sim \mu^{-1}n$. Therefore, if $\mu < \frac{\lambda}{1+\lambda}$, the probability that $a_k < x_k$ for all k is positive. This event, in turn, implies success of the algorithm. \square

5 Higher-dimensional arguments

Let $X^z = (X_n^z)_{n=0,1,2,\dots}$ denote a random walk starting at $z \in \mathbb{Z}^2$ and jumping according to $p(\cdot)$, the superindex is omitted when $z = \mathbf{0}$.

[She10]. $\mu_c \leq 1$.

Proof. Let V_n be the discrete ball of radius n around $\mathbf{0}$. Label and order all the particles initially found in V_n . We move the first particle until it is alone in a site, or until it exits the box. This can be done by subsequently toppling the site that contains the particle until one of these two conditions are met. If the particle is initially alone, there is nothing to do. After we are done with the first particle, we move the second particle until it is alone or exits the box. Then we move the third particle, and so on.

The path performed by each particle in this procedure is distributed as a random walk truncated at the time when it reached an empty site (we do not care about what the particle does once it exits V_n). In order to complete the path, we sample a random walk starting from the point where the particle was stopped. This second part of the path, which is not actually performed by the particle, is called a *ghost*. See Figure 10.

Let W count the *total* number of walks that visit $\mathbf{0}$ before leaving V_n . Let L count the number of walks that visit $\mathbf{0}$ before leaving V_n , but which do so *as ghosts*. We are interested in counting the number of such walks that visit $\mathbf{0}$ *as particles*, which is given by $W - L$.

It may be complicated to determine the set of locations where ghosts start. But we can

use the fact that at most one ghost starts at each site of V_n . Let us start an *artificial ghost* from each site where no actual ghost has been started, as also illustrated in Figure 10, so that exactly one ghost (artificial or not) starts from each site. Let $\tilde{L} \geq L$ denote total number of ghosts which visit $\mathbf{0}$ before leaving V_n .




Figure 10: Particles which visit $\mathbf{0}$ before or after exiting V_n , ghosts started each time a particle find itself alone at a site, and artificial ghosts started from other sites.

Although W and \tilde{L} are far from being independent, we claim that each of them has variance bounded by a multiple of its expectation, and thus is concentrated around its mean. Finally, we also claim that

$$\mathbb{E}[W] = \mu \mathbb{E}[\tilde{L}] \rightarrow \infty \text{ as } n \rightarrow \infty.$$

Hence, for $\mu > 1$, the weak law of large numbers gives $W - L \rightarrow \infty$ in probability as $n \rightarrow \infty$, proving non-fixation.

To complete the proof, let us prove the above elementary claims. First,

$$\mathbb{E}[\tilde{L}] = \sum_{x \in V_n} q_x, \quad \text{where} \quad q_x = \mathbb{P}[X^x \text{ visits } \mathbf{0} \text{ before exiting } V_n].$$

Also, $\mathbb{E}[W|\eta_0] = \sum_x \eta_0(x) q_x$, so $\mathbb{E}[W] = \mathbb{E}\mathbb{E}[W|\eta_0] = \mu \mathbb{E}[\tilde{L}]$. Write $m = \lceil n/3 \rceil$. Estim-

ing the above sum gives

$$\begin{aligned}
\mathbb{E}[\tilde{L}] &\geq \sum_{x \in V_m} \mathbb{P}[X^x \text{ visits } \mathbf{0} \text{ before exiting } V_m + x] \\
&= \sum_{z \in V_m} \mathbb{P}[X^{\mathbf{0}} \text{ visits } z \text{ before exiting } V_m] \\
&= \mathbb{E}[\text{number of sites visited by } X^{\mathbf{0}} \text{ before exiting } V_m] \rightarrow \infty
\end{aligned}$$

as $n \rightarrow \infty$. For the variance, first notice that \tilde{L} is given by the sum of independent indicator functions, so $\mathbb{V}\tilde{L} \leq \mathbb{E}\tilde{L}$. Conditioned on η_0 , W is also given by the sum of independent indicator functions, so $\mathbb{E}\mathbb{V}[W|\eta_0] \leq \mathbb{E}\mathbb{E}[W|\eta_0] = \mathbb{E}W$ and $\mathbb{V}\mathbb{E}[W|\eta_0] = \sum_x \mathbb{V}[\eta_0(x)]q_x = \mu\mu^{-1}\mathbb{E}[W]$, where μ denotes first the variance and then the expectation of a Poisson. Finally, the law of total variance gives $\mathbb{V}[W] \leq 2\mathbb{E}[W]$, finishing the proof. \square

[Tag14] (Biased walks). If the the walk is sufficiently directed, then there exist $\lambda > 0$ and $\mu < 1$ for which there is no fixation.

Proof. Denote the drift by $\mathbf{v} = \sum_z p(z)z \in \mathbb{R}^d \setminus \{\mathbf{0}\}$. We assume that $\mathbf{v} \cdot \mathbf{e}_1 > 0$ by symmetry, and that $d = 2$ for simplicity. We say that the walk X^z is *good* if $(X_n^z - z) \cdot \mathbf{e}_1 > 0$ for all $n \geq 1$. Let

$$\mathcal{K} = \mathbb{P}[(X_n)_n \text{ is good}] > 0.$$

Notice that \mathcal{K} tends to 0 as the walk becomes less biased, and tends to 1 when the probability of jumping to the left is small. We will show that the system stays active if

$$\left[\mu - \frac{\lambda}{1+\lambda}(1 - e^{-\mu})\right]\mathcal{K} > e^{-\mu}.$$

This condition is meaningful when $\mathcal{K} > \frac{1}{e}$.¹ The result is weaker than in $d = 1$ because the argument requires some control on the position where particles exit a large box.

Let $V_n = \{z = (x, y) \in \mathbb{Z}^2 : x \in [-n, 1], y \in [-n^3, n^3]\}$. Order these sites by writing $V_n = \{z_1, z_2, \dots, z_N\}$, where $N = 2n^4$, observing that sites with smaller x must appear first. We now stabilize V_n as follows.

Pick one particle initially present in z_1 , and move that particle until it either sleeps or exits V_n . If it sleeps in a given site of V_n , we leave that particle in that site and do not touch it anymore. In this case we start a *ghost* from that site so that the whole path associated to that particle is distributed as a random walk, part of which is performed by the particle,

¹For Bernoulli initial distribution, the result is a bit stronger. The term $e^{-\mu}$ is replaced by $1 - \mu$. In this case, the above condition gives $\mu_c < 1$ for any $\lambda < \infty$ and any $\mathcal{K} > 0$.

and part of it by the ghost. We say that the particle is good if the corresponding walk is good.

Repeat the same procedure for all particles initially present in z_1 , except for the last particle. When a single particle is left at z_1 , its first action may be to sleep, otherwise it is to jump. If it jumps we treat it as the previous particles, so that we assign a walk to it (part of which may consist of a ghost walk). If it sleeps, we declare that particle *dead*, and do not even assign it a walk.

After we are finished with all particles initially present in z_1 , we repeat the same action with particles found on z_2 , then z_3 , and so on.

Sites which contain a dead particle are those initially containing at least one particle and such that the last particle sleeps instead of performing its first jump, so the density of dead particles is $(1 - e^{-\mu}) \frac{\lambda}{1+\lambda}$. The total number of walks is given by the number of particles minus the number of dead particles, so it has density

$$\tilde{\mu} := \mu - \frac{\lambda}{1+\lambda}(1 - e^{-\mu}).$$

When the particle is good, its walk is always to the right of its initial position, and the particles initially present on this region have not yet been moved. Therefore, a good particle can only sleep at the set of sites that were initially empty, which has density $e^{-\mu}$.

These observations explain the requirement $\tilde{\mu}\mathcal{K} > e^{-\mu}$. Let us start a ghost walk at every site that was initially empty, even if no particle was actually stopped there. This way we obtain an upper bound for the set of ghosts associated to good particles. It thus follows by the LLN that at least $(\tilde{\mu}\mathcal{K} - e^{-\mu})N - o(N)$ good particles exit V_n before sleeping. However, we want to estimate the number of particles that exit V_n through site $\mathbf{0}$.

Let W count the number of good walks that exit V_n through $\mathbf{0}$ in this procedure. Let L count the number of good walks that exit V_n through $\mathbf{0}$, but which do so as ghosts. We are interested in the number $W - L$ of good particles that exit V_n through $\mathbf{0}$ before sleeping. In principle it is hard to estimate L , so we consider the number $\tilde{L} \geq L$ of ghosts that exit V_n through $\mathbf{0}$, when ghosts are started from each $z \in V_n$ that was initially empty.

Since both W and \tilde{L} are given by sums of independent variables whose variance is comparable to the expectation (see the proof of [She10] above), it suffices to show that

$$\mathbb{E}[W] \geq \tilde{\mu}\mathcal{K}n - o(n) \quad \text{and} \quad \mathbb{E}[\tilde{L}] \leq e^{-\mu}n$$

in order to have $W - \tilde{L} \rightarrow \infty$ as $n \rightarrow \infty$, in probability, which in turn implies non-fixation.

We use a simple argument of symmetry. First, writing $G_k = \{(x, y) \in \mathbb{Z}^2 : x = k\}$,

$$\begin{aligned}\mathbb{E}[\tilde{L}] &= \sum_{x=-n}^{-1} \sum_{y=-n^3}^{n^3} e^{-\mu} \mathbb{P}[X^{x,y} \text{ exits } V_n \text{ through } \mathbf{0}] \\ &\leq e^{-\mu} \sum_{x=-n}^{-1} \sum_{y \in \mathbb{Z}} \mathbb{P}[X^{x,y} \text{ reaches } G_0 \text{ at } \mathbf{0}] \\ &= e^{-\mu} \sum_{x=-n}^{-1} \sum_{y \in \mathbb{Z}} \mathbb{P}[X^{x,0} \text{ reaches } G_0 \text{ at } (0, -y)] = e^{-\mu} n.\end{aligned}$$

Finally, writing $H_k = \{(x, y) \in \mathbb{Z}^2 : y = \pm k\}$,

$$\begin{aligned}\mathbb{E}[W] &= \sum_{x=-n}^{-1} \sum_{y=-n^3}^{n^3} \tilde{\mu} \mathbb{P}[X^{x,y} \text{ is good and exits } V_n \text{ through } \mathbf{0}] \\ &\geq \tilde{\mu} \sum_{x=-n}^{-1} \sum_{y=-n^2}^{n^2} \mathbb{P}[X^{x,y} \text{ is good and reaches } G_0 \text{ at } \mathbf{0}, \text{ before reaching } y + H_{n^2}] \\ &= \tilde{\mu} \sum_{x=-n}^{-1} \sum_{y=-n^2}^{n^2} \mathbb{P}[X^{x,0} \text{ is good and reaches } G_0 \text{ at } (0, -y), \text{ before reaching } H_{n^2}] \\ &= \tilde{\mu} \sum_{x=-n}^{-1} \mathbb{P}[X^{x,0} \text{ is good and reaches } G_0 \text{ before } H_{n^2}] \\ &\geq \tilde{\mu} \sum_{x=-n}^{-1} \mathbb{P}[X^{x,0} \text{ is good and reaches } G_{x+n} \text{ before } H_{n^2}] \\ &= \tilde{\mu} n \mathbb{P}[X \text{ is good and reaches } G_n \text{ before } H_{n^2}].\end{aligned}$$

Since $\frac{X_k}{k} \rightarrow \mathbf{v}$ as $k \rightarrow \infty$, the last probability tends to \mathcal{K} as $n \rightarrow \infty$. This proves the desired inequalities, thus finishing the proof. \square

[She10]. For $\lambda = \infty$, $\mu_c > 0$.

Proof. We will in fact show that, if μ is small, then for almost every initial configuration η_0 , for any realization of \mathcal{I} , only finitely many particles can visit the origin. The idea is to assume that an adversary is trying to bring as many particles to the origin as possible, being able to move the particles to any nearest neighbor rather than following the random jumps. The proof uses the percolative structure of low-density sets.

Given η_0 , for each finite set $V \subseteq \mathbb{Z}^d$, define its *weight* by $w(V) = \sum_{x \in V} \eta_0(x)$. A finite connected set $V \subseteq \mathbb{Z}^d$ is *internally fillable* if $w(V) \geq |V|$.

Consider the decomposition of \mathbb{Z}^d into clusters \mathcal{C} given by

$$\mathcal{C}(x) = \bigcup \left\{ V \subseteq \mathbb{Z}^d : V \text{ is finite and connected, } x \in V, \text{ and } w(V) \geq |V| \text{ or } |V| = 1 \right\}.$$

These are called clusters because $\mathcal{C}(x) = \mathcal{C}(y)$ whenever $y \in \mathcal{C}(x)$.

The main observation is that, regardless of \mathcal{I} , for a particle starting at y to visit site x , it is necessary that some finite connected set V containing both x and y is internally fillable. In other words, it is necessary that $y \in \mathcal{C}(x)$.

Therefore, in order to show fixation, it suffices to prove that, with positive probability, $|\mathcal{C}(x)| < \infty$. This condition is equivalent to

$$\sup \left\{ |V| : \frac{w(V)}{|V|} \geq 1 \right\} < \infty.$$

It follows from results in [Mar02] about “greedy lattice animals” that, for some $C > 0$,

$$\text{a.s. } \lim_{n \rightarrow \infty} \max_{V: |V|=n} \frac{w(V)}{|V|} \leq C \mathbb{E}[\eta_0(\mathbf{0})^{d+1}].$$

Since we are assuming Poisson initial distribution, the bound can be made less than 1 by choosing μ small enough, which finishes the proof. \square

6 A multi-scale argument

[ST14] (Unbiased walks). For any $\lambda > 0$, $\mu_c > 0$.

Overview of the proof. The proof is too long and technical for these notes. We give an overview of the general strategy, omitting many delicate points.

The main step is to show that an initial configuration restricted to a very large box stabilizes within a slightly larger box, with high probability. This is proved by recursion on the scale of the box, so the proof looks very little into the details of the actual ARW dynamics. In a sense, this kind of approach fits to our intuition that no matter how big a defect is, it will only affect a neighborhood of comparable size.

The box at scale k is a cube of length L_k , defined as follows. Let $\gamma = 1/10$, $L_0 = 10^4$ and

$$L_{k+1} = \lfloor L_k^\gamma \rfloor^2 L_k.$$

Notice that L_k increases as a doubly exponential of k . We also let $R_{k+1} = \lfloor L_k^\gamma \rfloor L_k$ as an intermediate scale between L_k and L_{k+1} . In Figure 11 we see an *inner box* B'_{k+1} , an intermediate box and a *full box* B_{k+1} of level $k+1$.




Figure 11: Boxes and scales; $L_k \ll R_{k+1} \ll L_{k+1}$.

Let p_k denote the probability that, starting from a Poisson in B'_k , some particle exits B_k . That $p_k \rightarrow 0$ fast as $k \rightarrow \infty$ follows from the recursion relation

$$p_{k+1} \leq \frac{L_{k+1}^{2d}}{L_k^{2d}} p_k^2 + e_{k+1},$$

consisting of a combinatorial term, the probability p_k^2 that stabilization fails twice at scale k , and the probability e_{k+1} that something goes wrong at scale $k+1$. Indeed, if p_{k_0} is small enough and $e_k \rightarrow 0$ fast enough, then the square power above beats the $1 + 2\gamma$ power in the definition of L_k , and p_k vanishes doubly-exponentially fast in k .

Let us describe some aspects of this recursion step, depicted in Figure 12.

States in light gray have Poisson product distribution with the right density. They are restricted to the inner box of level $k+1$ for the initial configuration, and the inner boxes of level k for the “sieved configurations”. States in dark gray are *stable states*, typically attained by the dynamics. States in gray are “balanced configurations”. Thick arrows




Figure 12: Illustrative diagram of events for the recursion relation

represent typical events, while thin arrows represent events of low probability, either e_{k+1} or p_k .

Starting fresh. To let the dynamics run on boxes of the previous scale and use recursion, it is important to start with a Poisson product distribution within their inner boxes. This is achieved by a sieving procedure described below.

Worst case scenario. If the dynamics fail to stabilize all the $\frac{L_{k+1}^d}{L_k^d}$ boxes of level k , the configuration inside these boxes is no longer Poisson. In the absence of any useful knowledge about the resulting distribution of particles in this case, we use only the fact that the total number of particles within each box is still a Poisson random variable, and thus cannot be much larger than its mean. A *balanced configuration* is such that the number of particles within each box of level k is appropriately bounded.

Sieving procedure. Starting from a balanced configurations, we move each particle for quite enough time until their relative position within the box of level k that contains it has mixed. If the particle happens not to be in its inner box, we repeat the procedure until it is. This *reshuffling* with *sieving* results in a state that can be coupled with an i.i.d. Poisson configuration with high probability. This is one of the heaviest statements in [ST14] and not simple to prove. In order for this coupling to be possible, a slight increase in the density is necessary, analogous to the *sprinkling* technique in percolation (this increase should decay just fast enough so that it is summable over k).

The chain of events. By hypothesis we start with a Poisson product measure inside the inner box B'_{k+1} . The first good event is that such a configuration is balanced, that is, each box B_k has a bounded number of particles. We then let these particles remix by forcing them to jump. During this procedure they can reach the intermediate box, but not the full box. The second good event is that the resulting configuration is properly sieved. We now let the evolution run normally within each box B_k , and the third good event is that each box stabilizes nicely without letting particles leave. It may happen that some of these boxes of level k is not stabilized. In this case, the configuration is still balanced with high probability. We remix them again, now obtaining a sieved configuration in the full box B_{k+1} . The system is given a second chance to stabilize, which should happen typically, but may fail again if again some of the level- k boxes does not stabilize as expected.

In the proof there are several aspects to keep under control, and several delicate statements that we omit here. The above description is not intended to let the reader understand the sketch of proof, but hopefully gives a general flavor of the main argument. For all the details the reader is referred to the original article. \square

7 Arguments using particle-wise constructions

The techniques presented so far used the Diaconis-Fulton construction and its properties, as described in [RS12]. In the particle-wise construction, the randomness of the jumps is not attached to the sites, but to the particles. In this section we consider techniques which use the particle-wise construction, or sometimes a combination of both.

7.1 Preliminaries

System with labeled particles Each existing particle at time $t = 0$ is assigned a label (x, j) , where $x \in \mathbb{Z}^d$ denotes its starting position and $j = 0, \dots, \eta_0(x) - 1$ distinguish particles starting at the same site x . Let $Y^{x,j} = (Y_t^{x,j})_{t \geq 0}$ be given by the position of particle (x, j) at each time t . We write $Y_t^{x,j} = \Upsilon \notin \mathbb{Z}^d$ if $\eta_0(x) \geq j$. Let $\gamma^{x,j} = (\gamma^{x,j}(t))_{t \geq 0}$ be given by $\gamma^{x,j}(t) = 1$ if particle (x, j) is active at time t or ϱ if it is passive. Write $\mathbf{Y} = (Y^{x,j})_{x,j}$ and $\boldsymbol{\gamma} = (\gamma^{x,j})_{x,j}$. Then the triple $\boldsymbol{\zeta} = (\eta_0, \mathbf{Y}, \boldsymbol{\gamma})$ describes the whole evolution of the system.

Whereas the process $(\eta_t)_{t \geq 0}$, given by

$$\eta_t(z) = \sum_x \sum_{j < \eta_0(x)} \gamma^{x,j}(t) \cdot \delta_{Y^{x,j}(t)}(z),$$

only counts the number of particles at a given site at a given time, having each particle labeled gives a lot more information and allows different techniques to be employed.

In a system whose initial configuration η_0 contains finitely many particles, the evolution described above is always well defined, since it is simply a continuous-time Markov chain on a countable state-space. Many different constructions will produce $\boldsymbol{\zeta}$ with the correct distribution.

Particle-wise randomness Assign to each particle (x, j) a continuous-time walk $X^{x,j} = (X_t^{x,j})_{t \geq 0}$, independently of anything else, as well as a Poisson clock $\mathcal{P}^{x,j} \subseteq \mathbb{R}_+$ according to which the particle will try to sleep. $X^{x,j}$ is the path of the particle parameterized by its *inner time*, which may be slowed down with respect to the system time, depending on the interaction with other particles. (In the ARW model, when the particle is passive, the inner time halts until it is re-activated by another particle.) For this reason, $X^{x,j}$ will be called the *putative trajectory* of particle (x, j) . These random elements will be denoted by $\boldsymbol{\xi} = (\eta_0, \mathbf{X}, \boldsymbol{\mathcal{P}})$. Let $\sigma^{x,j}(t)$ denote the inner time of particle (x, j) at instant t . Then $(Y_t^{x,j})_{t \geq 0}$ is given by $Y_t^{x,j} = X_{\sigma^{x,j}(t)}^{x,j}$.

Construction of the infinite system When η_0 contains finitely many particles only, ζ is determined by ξ in the obvious way, and it works for a.e. \mathbf{X} and \mathcal{P} . The simultaneous construction on the whole space, is done on the sequences of balls $B(y, n)$, $n \in \mathbb{N}$, centered at each site $y \in \mathbb{Z}^d$. This family of sequences is countable and translation-invariant. For each ball, consider the evolution of the system with initial configuration $\eta_0 \cdot \mathbb{1}_{B(y, n)}$. We say that this construction is *well defined* if (i) for each x and j , and for each $T > 0$, both $Y_{[0, T]}^{x, j}$ and $\gamma_{[0, T]}^{x, j}$ are the same in the systems $(\eta_0 \cdot \mathbb{1}_{B(y, n)}, \mathbf{Y}, \gamma)$ for all but finitely many n , and (ii) the limiting process does not depend on y .

The above definition is convenient for two reasons. First, the family of finite approximations to \mathbb{Z}^d is countable, which makes it possible to prove that the construction is a.s. well defined. Second, this family is translation invariant, so that, whenever the construction is well defined, it is translation covariant. More precisely, for any translation θ of \mathbb{Z}^d , if $\zeta(\xi)$ is well defined, then $\zeta(\theta\xi)$ is well defined and equals $\theta\zeta(\xi)$, where $\theta\xi = (\theta\eta_0, \theta\mathbf{X}, \theta\mathcal{P})$, etc. In particular, this implies that the system is ergodic, satisfies the mass transport principle, and can be approximated by finite systems regardless of which construction is used.

The Mass-Transport Principle Let $m : \mathbb{Z}^d \times \mathbb{Z}^d \rightarrow \mathbb{R}_+$ be a translation-invariant random function, that is, $m(x, y; \xi) = m(\theta x, \theta y; \theta\xi)$ for any translation θ . The Mass Transport Principle is given by

$$\mathbb{E}\left[\sum_y m(x, y)\right] = \mathbb{E}\left[\sum_y m(y, x)\right].$$

Informally, the MTP says that the amount of mass transmitted from a vertex x is equal to the amount of mass entering x . It seems like an innocent and perhaps obvious identity, but its strength lies in its versatility, since it holds for *any* such function. The proof consists simply on re-indexing the sum and using translation invariance. See [LP, Chapter 8] for applications and generalizations to other settings.

The particle-hole model Let us introduce a related model that will be useful in the sequel. Particles perform continuous-time random walks independently of each other. Sites not containing any particle are called *holes*. When a particle is alone at some site, it *settles* there forever, filling the corresponding hole. After the hole has been filled, the site becomes available for other particles to go through. If a site is occupied by several particles at $t = 0^-$, we choose one of them uniformly to fill the hole at $t = 0$, and the other particles remain free to move.

As in the ARW with $\lambda = \infty$, once a site has at least one particle, it will always retain one particle. The differences are (i) sites with n particles are toppled at rate n instead of $n - 1$ and (ii) in the system with labeled particles, it is the first particle to arrive at a site that is retained, whereas in the ARW the particles can replace each other. Nevertheless, both models have the same Diaconis-Fulton representation, and therefore they are equivalent in terms of fixation at $x = \mathbf{0}$. The Poisson clocks \mathcal{P} are not used in either system.

7.2 Results

Fixation as defined so far concerns the state of sites, and will be called *site fixation*. When each labeled particle eventually fixates we call that *particle fixation*.

[AGG10] Assume that the particle-wise construction is a.s. well defined. If some particles are not fixating, then sites are not fixating. Therefore, $\mu_c \leq 1$.

Let us see how $\mu_c \leq 1$ follows from the main claim. First, it follows from the MTP that:

The density of particles that fixate cannot be larger than 1.

Indeed, let $A(x, j, y)$ denote the event that particle (x, j) fixates at site y . Let $m(x, y) = \sum_j \mathbf{1}_{A(x, j, y)}$. The density of particles that fixate is given by $\mathbb{E} \sum_y m(x, y)$ and the density of sites eventually occupied by a fixating particle is given by $\mathbb{E} \sum_y m(y, x)$. These are equal by the MTP. Since at most one particle can fixate at a given site, $\sum_y m(y, x) \leq 1$, proving the estimate. Now, if $\mu > 1$, there is a positive density of particles which do not fixate, which implies that there is no site fixation. Therefore, $\mu_c \leq 1$.

Proof. Let $A^{v, j}$ denote the event “ $\eta_0(v) > j$ and (v, j) is non-fixating” and write $A^v = A^{v, 0}$. Suppose the probability that some particle is non-fixating is positive. Then there is j such that $\mathbb{P}(A^{0, j}) > 0$. By interchangeability of particles, $a := \mathbb{P}(A^0) = \mathbb{P}(\eta_0(\mathbf{0}) > 0, (\mathbf{0}, 0) \text{ non-fixating}) \geq \mathbb{P}(\eta_0(\mathbf{0}) > j, (\mathbf{0}, 0) \text{ non-fixating}) = \mathbb{P}(A^{0, j}) > 0$.

As a warm up, notice that by the MTP the number M_t of non-fixating particles present at $\mathbf{0}$ at time t satisfies $\mathbb{E} M_t \geq a$, for any $t > 0$. Hence, the $\liminf_t \mathbb{E} M_t \geq a$. If we had some control on $\mathbb{V} M_t$, we could conclude that $\liminf_t \mathbb{P}(M_t \geq 1) > 0$, implying non-fixation. To control variance, we will use a local condition instead, plus extra randomization.

Since the system ζ is a measurable function of the randomness ξ , for any $\epsilon > 0$ there is $k \in \mathbb{N}$ such that the event A^0 can be ϵ -approximated by some event B^0 that depends only on $(\eta_0(x), X^x, \mathcal{P}^x)$ for $\|x\| \leq k$. Let B^v denote the corresponding translation of the event

B^0 . When B^v occurs, we say that particle $(v, 0)$ is a *candidate*. It is a *good candidate* if A^v also occurs, otherwise it is a *bad candidate*.

Fix $T > 0$. For $n \in \mathbb{N}$, let us add more randomness to the system by sampling Z^v uniformly amongst the first n different sites in the putative trajectory $X^{v,0}$ after time T , independently over v . Define $C(v, u)$ as the event “ B^v and $Z^v = u$ ”. Let

$$q(v, u) = \mathbb{P}(C(v, u) \mid \xi) \quad \text{and} \quad Q(u) = \sum_v q(v, u).$$

By the mass-transport principle,

$$\mathbb{E}[Q(v)] = \sum_u \mathbb{P}(C(u, v)) = \sum_u \mathbb{P}(C(v, u)) = \mathbb{P}(B^v) =: b > 0.$$

Notice that $q(v, u) \leq \frac{1}{n}$. Notice also that $q(v, u)$ and $q(w, u)$ are independent if $|w - u| > 2k$. Using these two facts, it can be shown that, as $n \rightarrow \infty$, $\mathbb{V}[Q(v)] \rightarrow 0$, and thus $Q(v) \rightarrow b$ in probability.

Let $N(v) = \sum_u \mathbb{1}_{C(u, v)}$ count the number of candidates for which $Z^u = v$. Then

$$\mathbb{P}(N(v) = 0 \mid \xi) = \prod_u (1 - q(u, v)) \leq e^{-Q(v)} \rightarrow e^{-b} \text{ in probability as } n \rightarrow \infty.$$

Also, let $\tilde{N}(v) = \sum_u \mathbb{1}_{C(u, v) \setminus A^u}$ count the number of bad candidates for which $Z^u = v$. Then, using the mass-transport principle,

$$\mathbb{E}[\tilde{N}(v)] = \sum_u \mathbb{P}[C(u, v) \setminus A^u] = \sum_u \mathbb{P}[C(v, u) \setminus A^v] = \mathbb{P}[B^v \setminus A^v] \leq \epsilon.$$

Let D^v denote the event “there exists a good candidate $(u, 0)$ such that $Z^u = v$ ”. Using the two last estimates we get

$$\mathbb{P}(D^v) \geq \mathbb{P}(N(v) \geq 1) - \mathbb{P}(\tilde{N}(v) \geq 1) \geq 1 - e^{-a+\epsilon} - \delta_n - \epsilon,$$

where $\delta_n \rightarrow 0$ as $n \rightarrow \infty$. Choosing ϵ small and n large, we have $\mathbb{P}(D^0) > \frac{a}{2}$.

On the event D^0 , there is some non-fixating particle $(u, 0)$ and some inner time $t > T$ such that $X_t^{u,0} = \mathbf{0}$, so vertex $\mathbf{0}$ is visited by an active particle after time T . Letting $T \rightarrow \infty$, we get $\mathbb{P}(\text{site } \mathbf{0} \text{ not fixating}) \geq \frac{a}{2}$, and by the 0-1 law $\mathbb{P}(\mathbf{0} \text{ not fixating}) = 1$. \square

[CRS14] For $\lambda = \infty$, $\mu_c \geq 1$.

We start with the following observation about the particle-hole model:

The density of holes filled by time t equals the density of particles settled by time t .

To see that, let $A(x, j, y)$ denote the event that particle (x, j) settles at site y by time t , and let $m(x, y) = \sum_j \mathbf{1}_{A(x, j, y)}$. The density of particles settled by time t is given by $\mathbb{E} \sum_y m(x, y)$ and the density of holes filled by time t is given by $\mathbb{E} \sum_y m(y, x)$. These are equal by the MTP, proving the identity.

Proof. Suppose that $\mu < 1$. Using the above identity,

$$\mathbb{P}[\mathbf{o} \text{ contains an unfilled hole at time } t] \geq 1 - \mu > 0.$$

Since the above event is decreasing in t and the lower bound is uniform with respect to t ,

$$\mathbb{P}[\mathbf{o} \text{ is never visited}] > 0.$$

In particular, the probability that $\mathbf{0}$ is visited finitely many times in the particle-hole model is positive, so by the 0-1 law, $\mathbb{P}[m_{\eta_0}(\mathbf{0}) < \infty] = 1$, which means that the ARW with $\lambda = \infty$ fixates. \square

[CRS14] No fixation when $\mu = 1$.

The proof is done in full details in [CRS13] for two-type annihilating systems, that is, with reaction $A + B \rightarrow \emptyset$. Active particles are particles of type A , and holes are particles of type B . The reaction happens when a particle settles, thereby filling a hole.

Sketch of the proof. Suppose that sites fixate in the particle-hole model. By the first result presented in this section, site fixation implies that every particle eventually settles. Letting $t \rightarrow \infty$ in the previous observation,

$$\mathbb{P}[\mathbf{0} \text{ is ever visited}] = \mu.$$

In the sequel we will show that, under the assumption of site fixation,

$$\mathbb{P}[\mathbf{0} \text{ is never visited}] > 0.$$

This implies that $\mu < 1$, therefore proving that there cannot be fixation for $\mu = 1$.

Assuming site fixation, necessarily, there exists $k \in \mathbb{N}$ such that

$$\mathbb{P}[\text{the number of particles which ever visit } \mathbf{0} \text{ equals } k] > 0.$$

Moreover, there exist $x_1, \dots, x_k \in \mathbb{Z}^d$ such that $\mathbb{P}[\mathcal{U}] > 0$, where

$$\mathcal{U} = [\text{the particles which ever visit } \mathbf{0} \text{ are initially at the sites } x_1, \dots, x_k].$$

Consider two copies ξ and $\tilde{\xi}$ of the system, coupled as follows. We take $\tilde{\mathbf{X}} = \mathbf{X}$, and $\tilde{\eta}_0(x) = \eta_0(x)$ for $x \notin \{x_1, \dots, x_k\}$. Finally, for $x \in \{x_1, \dots, x_k\}$, we sample $\tilde{\eta}_0$ and η_0 independently. Now notice that

$$\begin{aligned} \mathbb{P}[\mathcal{U} \text{ occurs for } \tilde{\xi}, \text{ and } \eta_0(x_1) = \dots = \eta_0(x_k) = 0] &= \\ &= \mathbb{P}[\mathcal{U} \text{ occurs for } \tilde{\xi}] \times \mathbb{P}[\eta_0(x_1) = \dots = \eta_0(x_k) = 0 \text{ for } \xi] > 0. \end{aligned}$$

On the above event, no particle ever visits $\mathbf{0}$ in the system ξ . Indeed, on the above event, the initial configuration of ξ is the same as that of $\tilde{\xi}$ except for the deletion of the particles present in $\{x_1, \dots, x_k\}$. In particular, all the particles which visit the origin in $\tilde{\xi}$ are deleted in ξ . Recalling that ξ and $\tilde{\xi}$ share the same putative trajectories, by following how the effect of deleting such particles propagates in the system evolution, we see that in the system ξ no particles can ever visit $\mathbf{0}$. \square

References

- [AGG10] G. AMIR, O. GUREL-GUREVICH. *On fixation of activated random walks*. Elect Comm in Probab **15** (2010), pp. 119–123.
- [CRS13] M. CABEZAS, L. T. ROLLA, V. SIDORAVICIUS. *Recurrence and density decay for diffusion-limited annihilating systems*, 2013. Preprint. arXiv:1309.4387.
- [CRS14] ———. *Non-equilibrium phase transitions: Activated random walks at criticality*. J Stat Phys **155** (2014), pp. 1112–1125.
- [DRS10] R. DICKMAN, L. T. ROLLA, V. SIDORAVICIUS. *Activated random walkers: Facts, conjectures and challenges*. J Stat Phys **138** (2010), pp. 126–142.
- [KS06] H. KESTEN, V. SIDORAVICIUS. *A phase transition in a model for the spread of an infection*. Illinois J Math **50** (2006), pp. 547–634.
- [LP] R. LYONS, Y. PERES. *Probability on trees and networks*. Book in preparation. Current version available at <http://mypage.iu.edu/~rdlyons>.

- [Mar02] J. B. MARTIN. *Linear growth for greedy lattice animals*. Stochastic Process Appl **98** (2002), pp. 43–66.
- [RS12] L. T. ROLLA, V. SIDORAVICIUS. *Absorbing-state phase transition for driven-dissipative stochastic dynamics on Z* . Invent Math **188** (2012), pp. 127–150. ArXiv:0908.1152.
- [She10] E. SHELLEF. *Nonfixation for activated random walks*. Alea **7** (2010), pp. 137–149.
- [ST14] V. SIDORAVICIUS, A. TEIXEIRA. *Absorbing-state transition for stochastic sand-piles and activated random walks*, 2014. Preprint. arXiv:1412.7098.
- [Tag14] L. TAGGI. *Active phase for activated random walk with density below unit*, 2014. Preprint. arXiv:1403.1986.